ERCAP Request #83085 for FY2009

1. Principal Investigator

Name	Phone	Email
Bethel, Edward (Wes)	510-486-7353	ewbethel@lbl.gov

2. Other Authorized Preparers

3. Senior Investigators	(Utah) Chris Johnson, Chuck Hansen, Valerio Pascucci, Claudio Silva (LLNL) Hank Childs, Brad Whitlock, Dan Laney (ORNL) Sean Ahern, Jeremy Meredith, Dave Pugmire, George Oustrochov (UCD) Ken Joy, Christoph Garth, Eduard Deines
4. Project Title	SciDAC2 Visualization and Analytics Center for Enabling Technologies
5. Project Name	VACET
6. Project Class	SciDAC
7. Sponsoring Site	USA: Lawrence Berkeley National Laboratory
8. Science Category	Computer Sciences
9. DOE Office and Program	ASCR - Computer Sciences

10. Is this project funded by the DOE Office of Science?

X

Yes, this project has direct grant support from the DOE Office of Science

Who is your DOE Program Manager?

Osni Marques and Fred Johnson

DOE Office of Science Grant Number(s)

1. SciDAC2 Visualization and Analytics Center for Enabling Technologies (VACET). Total award amount is $2.2M/yr for five years.
2. High Performance Visualization, ASCR/MICS Base Program. Award is $600K yr through FY11.

According to LBL Finance people, there are no "Grant

Numbers" associated with either of these awards but both are "on the books."

11.1 Does this project make use of national security information?

Yes X No

11.2 Does this project collect and store Social Security Numbers, Personally Identifiable Health Information (names or other identifiers matched with health data), Driver's License Numbers, or Financial Account Numbers?

Yes X No

11.3 If this project is given a NERSC award, I agree to monitor the usage associated with it to ensure that, to the best of my ability to determine, usage is for the project described here.

X Yes No

11.4 For continuing projects: I have audited the MPP and/or HPSS usage associated with this project, and to the best of my ability to determine, all usage was for the project specified.

X Yes No

12. Computational Resources Requested

Center	Resource	Alloc Type	Repo	Hours Used 2008	Hours Requested 2009
NERSC	MPP (MPP Charging)	DOE Production	m636	7,524	160,000

13. Mass Storage Resources Requested

Center	Resource	Alloc Type	Repo	SRUs Used 2008	SRUs Requested 2009
NERSC	HPSS Charging	DOE Production	m636	148,423	250,000

14. Justification for Resources Requested

For both the VACET and LBNL Base Program efforts, we anticipate performing several different types of activities to meet our project's objectives. The hours/storage we request reflect our best estimate of the resources needed to accomplish work in these areas:

1. Production visual data analysis: consists of making movies and images for stakeholders.
2. Visual data analysis research and development: we have multiple projects in the areas of: topologically based, temporal feature detection and tracking that will be applied to stakeholder projects in climate and combustion; visual data analysis and scientific data management (including machine learning, clustering, PCA/ICA) of
large, particle-based datasets for stakeholder projects in fusion and accelerator modeling; high performance parallel I/O projects.

During AY08, we requested and were awarded 50K hours, but ended up using about 150K hours. The overrun was the result of starting to work on parallel I/O projects for our SciDAC stakeholders. Last year, we requested 250K SRUs and have used about half of them. We have 1 TB quota on NGF, and NERSC has kindly accommodate temporary quota increases on an as-needed basis. We anticipate peak requirements this year for NGF to be in the 10TB-20TB range; this figure will only go up in the years ahead.

More information about the projects/request:

- Topological Analysis of AMR combustion and astrophysics simulation:
-- 20K hours for running analysis, preparing visual results.
-- 5TB of filesystem storage for source data, intermediate results. NGF is best for this kind of thing.

- High performance visual analysis of Accelerator data
-- 20K hours for analysis, vis, etc.
-- 10TB of filesystem storage for source data, intermediate results, final results.

- Climate Parallel I/O, Data Modeling and Multiresolution Visual Data Analysis
-- 120K hours for I/O benchmarks, debugging and tuning.
-- 4TB of filesystem storage for source data, intermediate and final results.

15. PDSF Repos

16. Project Description

16.1 Project Summary: Provide a brief project description. What will this project accomplish? What is the significance of this work?

Launched in 2006 as one of nine centers under the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC-2), VACET focuses on leveraging scientific visualization and analytics software technology as an enabling technology for increasing scientific productivity and insight. Advances in computational technology have resulted in an information big bang, which in turn has created a significant data understanding challenge. This challenge is widely acknowledged to be one of the primary bottlenecks in contemporary science. The vision for our Center is to respond directly to that challenge by adapting, extending, creating when necessary and deploying visualization and data understanding technologies for our science stakeholders at DOE's open computing facilities (NERSC/LBNL and NCCS/ORNL). Using an

organizational model as a Visualization and Analytics Center for Enabling Technologies (VACET), we are well positioned to be responsive to the needs of a diverse set of scientific stakeholders, including other SciDAC projects, in a coordinated fashion using a range of visualization, mathematics, statistics, computer and computational science and data management technologies.

More specifically, VACET will provide visualization and analytics software infrastructure to support the challenging data understanding needs of SciDAC Science Applications. This infrastructure will be deployed at DOE's open computing facilities, both NERSC/LBNL and NCCS/ORNL.

16.2 URL for a relevant web page.

http://www.vacet.org/

17. Accomplishments

Major Accomplishments

Topological analysis applied to combustion data (1): Working with J. Bell (CCSE, LBNL), our team has performed R&D resulting in the ability to provide a quantitative basis for the relationship between turbulence and combustion processes. Our approach is to use topological analysis to segment time-varying AMR combustion simulation data into "burning" and "non-burning regions," and to study the behavior and characteristics of these regions over time in the presence of varying levels of turbulence in a way not possible with traditional data analysis techniques. This accomplishment is a major new capability that will aid combustion researchers in designing more efficient engines, burners, and power plants.

Topological analysis applied to combustion data (2): Working with Jacqueline Chen (SNL-CA), our team has developed a new technique for performing quantitative analysis of "extinction pockets" in time-varying 3D data produced by S3D, a DNS combustion simulation. The idea is to track the evolution of these extinction pockets over time, and to study their characteristics in relationship to other simulation parameters. This new capability provides a quantitative basis for understanding complex phenomena in a way not possible with traditional visualization or analysis techniques.

Accelerator Modeling: Working with Cameron Geddes, our team has developed a new technique for unsupervised classification/analysis of data produced by Laser-Wakefield simulation data. This technique accurately identifies particles undergoing wakefield simulation using a combination of machine learning and graph-based algorithms: particles undergoing wakefield simulation have both high levels of acceleration as well as spatial coherence. This new technique will accelerate scientific knowledge discovery by replacing a manual search phase, where a scientist manually examines data from different timesteps to visually identify what set of conditions are associated with particle acceleration, with one that can run automatically.

Accelerator Modeling: Working with Cameron Geddes (LBNL), along with researchers from the SciDAC SDM Center and LBNL's Visualization Research program, our team has developed and tested new production-quality, parallel-capable software infrastructure that is suitable for performing visual data exploration on multi-terabyte datasets produced by laser wakefield simulation codes. This new capability, which can perform particle tracking on multi-TB datasets in a few seconds when run in parallel on Franklin, replaces a manual and serial processing step that formerly required hours of runtime.

Minor Accomplishments

In addition to the major ones listed above, having access to NERSC resources has enabled progress on a number of more minor accomplishments:

Remote Visualization Architectures. Our team developed a new remote visualization infrastructure that is suitable for use in remote delivery of imagery produced by parallel, hardware accelerated graphics and visualization applications.

Remote Visualization Architectures: Our team conducted a study whereby images from visualization applications are precomputed in tiled and multiresolution, stored on a central server, then made available to remote clients through a garden-variety webserver. This study evaluates system performance for different network conditions (bandwidth, latency), different image tiling parameters, and different parameters for a predictive image prefetch algorithm.

AMR Visualization studies: We used the resources at NERSC to perform large-scale data format conversion for use in several different research projects in query-driven visualization. One project focused on developing a new data structure and algorithm suitable for deployment using CUDA on GPUs. Another focused on the fundamental data structures and algorithms needed to perform query-driven visual data analysis on time-varying AMR datasets.

18. Relevant Publications: Enter in the order authors, title, journal

18.1 Refereed Publications: List all refereed publications in the last 12 months based on research using NERSC resources. You may include publications submitted to journals but not publications in preparation.

Daniela Ushizima, Oliver Ruebel, Prabhat, Gunther Weber, E. Wes Bethel, Cecilia Aragon, Cameron Geddes, Estelle Cormier-Michel, Bernd Hamann, Peter Messmer, Hans Hagen. "Automated Analysis for Detecting Beams in Laser Wakefield Simulations". 2008 Seventh International Conference on Machine Learning and Applications, Proceedings of IEEE ICMLA'08, 2008. LBNL-960E.

E. Wes Bethel, Hank Childs, Ajith Mascarenhas, Valerio Pascucci, and Prabhat. Scientific Data Managment Challenges in High Performance Visual Data Analysis. In Arie Shoshani and Doron Rotem, editors, Scientific Data Management: Challenges, Existing Technology, and Deployment. Chapman & Hall/CRC Press, 2008.

Luke J. Gosink, John C. Anderson, E. Wes Bethel, Kenneth I. Joy. "Query-Driven Visualization of Time-Varying Adaptive Mesh Refinement Data." IEEE Transactions on Visualization and Computer Graphics, 14(6), Proceedings of IEEE Visualization 2008 (To appear). LBNL-803E.

Luke J. Gosink, Kesheng Wu, E. Wes Bethel, John D. Owens, Kenneth I. Joy. "Bin-Hash Indexing: A Parallel Method for Fast Query Processing." Technical Report LBNL-729E, Lawrence Berkeley National Laboratory, Berkeley CA, USA, 94720, 2008.

O. Ruebel, Prabhat, K. Wu, H. Childs, J. Meredith, C.G.R. Geddes, E. Cormier-Michel, S. Ahern, G.H. Weber, P. Messmer, H. Hagen, B. Hamann and E.W. Bethel, "Application of High-performance Visual Analysis Methods to Laser Wakefield Particle Acceleration Data." Poster at IEEE Visualization 2008, Columbus, Ohio, October 19-24, 2008. (To appear) LBNL-952E.

O. Ruebel, Prabhat, K. Wu, H. Childs, J. Meredith, C.G.R. Geddes, E. Cormier-Michel, S. Ahern, G.H. Weber, P. Messmer, H. Hagen, B. Hamann and E.W. Bethel, "High Performance Multivariate Visual Data Exploration for Extremely Large Data." SC08, Austin TX, November, 2008

Sarah Poon, Rollin Thomas, Cecilia Aragon, and Brian Lee, "Context-Linked Virtual Assistants for Distributed Teams: An Astrophysics Case Study." CSCW 2008: ACM Conference on Computer Supported

Cooperative Work. (To appear, 2008.) Best Paper Award Nominee

Cecilia Aragon, Stephen Bailey, Sarah Poon, Karl Runge, and Rollin Thomas, "Sunfall: A Collaborative Visual Analytics System for Astrophysics." J. Phys.: Conf. Ser. 125 012091 (Proceedings of SciDAC 2008). LBNL-657E.

Cecilia Aragon, Sarah Poon, Gregory Aldering, Rollin Thomas, and Robert Quimby, "Using Visual Analytics to Maintain Situation Awareness in Astrophysics." Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (IEEE VAST 2008).

E. Wes Bethel, Chris Johnson, Charles Hansen, Claudio Silva, Steven Parker, Allen Sanderson, Lee Myers, Martin Cole, Xavier Tricoche, Sean Ahern, George Ostrouchov, Dave Pugmire, Jamison Daniel, Jeremy Meredith, Valerio Pascucci, Hank Childs, Peer-Timo Bremer, Ajith Mascarenhas, Ken Joy, Bernd Hamann, Christoph Garth, Cecilia Aragon, Gunther Weber, and Prabhat. "Seeing the Unseeable." SciDAC Review, Number 8, Summer 2008, pp 24-33. LBNL-472E.

O. Ruebel, G. H. Weber, M-Y Huang, E. W. Bethel, M. D. Biggin, C. C. Fowlkes, C. Luengo Hendriks, S. V. E. Ker??nen, M. Eisen, D. Knowles, J. Malik, H. Hagen and B. Hamann, "Integrating Data Clustering and Visualization for the Analysis of 3D Gene Expression Data." IEEE Transactions on Computational Biology and Bioinformatics. 2008. (To appear) LBNL-382E.

C. C. Fowlkes, C. L. Luengo Hendriks, S. V. E. Ker??nen, G. H. Weber, O. Ruebel, M.-Y. Huang, S. Chatoor, A. H. DePace, L. Simirenko, C. Henriquez, A. Beaton, R. Weiszmann, S. Celniker, B. Hamann, D. W. Knowles, M. D. Biggin, M. B. Eisen, J. Malik. "A Quantitative Spatio-temporal Atlas of Gene Expression in the Drosophila Blastoderm." Cell Vol 133, 364-374, 18 April 2008.

18.2 Other Publications: List up to 5 other relevant publications in the last 12 months, also based on research using NERSC resources.

B. Paul, S. Ahern, E. W. Bethel, E. Brugger, R. Cook, J. Daniel, K. Lewis, J. Owen, and D. Southard. Chromium Renderserver: Scalable and Open Remote Rendering Infrastructure. IEEE Transactions on Visualization and Computer Graphics, 14(3), May/June 2008. LBNL-63693.

J. Chen, I. Yoon and E. W. Bethel. "Interactive, Internet Delivery of Visualization via Structured, Prerendered Multiresolution Imagery." IEEE Transactions on Visualization and Computer Graphics, Volume 14, Number 2, March/April 2008. LBNL-62252.

19.1 Code and Application Descriptions

Code Name	VisIt
Description	Our project will run multiple codes. VisIt is a production-quality, parallel-capable visualization application the uses a client-server architecture for displaying results of visualization to a remote client. It has been used on the BG/L system at LLNL to perform visual data analysis of a grid containing 27 billion zones -- it has proven to be sufficiently scalable on all modern platforms. In addition, it has been ported to run on Jaguar, so we reasonably expect it will run on all platforms at NERSC. Custom topological analysis codes. This family

	of codes, which are experimental, are run using large, time-varying input. Each timestep of data is processed in serial; parallelization occurs by running many such single-timestep analyses in parallel. Limiting factors for these codes is single-node memory size and I/O bandwidth. For that reason, we expect most of this work to be conducted on davinci.
Mathematics	For the topological analysis work, please refer to the following Siggraph 2007 and IEEE Visualization 2007 papers authored by a VACET team member: V. Pascucci, G. Scorzelli, P.-T. Bremer, and A. Mascarenhas. "Robust On-line Computation of Reeb Graphs: Simplicity and Speed" In Proceeding of SIGGRAPH 2007 (ACM Transation on graphics), pp. 58-1 to 58-9, 2007. Attila Gyulassy, Vijay Natarajan, Bernd Hamann, Valerio Pascucci. "Efficient Computation of Morse-Smale Complexes for Three-dimensional Scalar Functions" In IEEE Transactions on Visualization and Computer Graphics, Proceedings of IEEE Visualization 2007, October 2007.
Numerical Techniques	
Machines	Franklin, davinci.
Percent of Allocation	20
Planned Processors	Ensembles of serial runs, each has a large memory footprint (8-64GB).
Number of Processors Description	Computational Requirements,Memory Required,Large number of serial jobs
Particle Methods	
Linear Solvers	
Eigensolvers	
Other Algorithms	
Languages	c,c++,UPC,Python,MPI,MPI/IO
Libraries	papi,hdf4,hdf5,netcdf,OpenGL, Xlib, Qt, Posix communication & I/O, Boxlib, Chombo

Performance Limits	disk space,MPI I/O,Posix I/O,Single Node Memory Performance,Need large memory footprint for those tasks that are serial
Performance Comments	For the most part, our applications are very data intensive. The balance of I/O to FLOPS is weighted much more in favor of I/O than FLOPS when compared to most "solver" type codes.
Checkpoint?	N
Code Plans	More exploration of using UPC for scalable visualization algorithms.

19.2 Code and Application Performance

20. Storage and Data Movement

20.1 How much space do you need?

Home (permanent) space (Gigabytes):	100
Scratch space (Terabytes):	20
NERSC Global Filesystem (Terabytes):	10
HPSS archival storage (Terabytes):	20

20.2 How do you move data? Between which systems?

To/From HPSS and NERSC systems: HSI
To/From HPSS and off-site systems: HSI (having current versions of HSI clients for common workstations is very helpful, thanks!)

When HSI fails to off-site systems (50% of the time, to use a rough estimate), then we're reduced to using SCP.

21. Data and I/O

21.1 Next year, how much data do you expect your project to produce at NERSC? (Terabyes)

1.0

21.2 Next year, how much data do you expect your project to move to NERSC from external sources? (Terabyes)

50

21.3 How do you share data? Between which systems?

For the most part, our projects use large-scale scientific data produced by other projects. The 50TB "move to NERSC" figure reflects an estimate of the aggregate size of data we'll use in AY09 that is provided by our science stakeholders. A majority, 75% (estimated), will already be resident at NERSC (computed at NERSC), while the remaining 25% would be moved to NERSC from external sites (e.g., ORNL).

The data "we generate" for the most part consists of images/movies and performance data.

21.4 How do you perform I/O? Do you write concurrently from all processors to shared files (parallel I/O) or one file per processor or node? Have you experienced problems or constraints due to I/O?

We do both collective and non-collective I/O. We have had lots of problems doing collective I/O on franklin, and are engaging with NERSC staff to investigate/resolve the problem.

21.5 Grid meshing techniques used

Grid Meshing (Q 21b)

no mesh,
structured rectilinear mesh,
structured curvilinear mesh,
unstructured mesh,
hierarchical adaptive mesh refinement,
hybrid,
2D,
3D,
4D,
gt4D,
10s of billions per timestep

21.6 What file formats / data libraries do you use?

File Formats (Q. 21.6)

POSIX,
MPI IO,
HDF5,
H5Part,
FITS,
FSML - fusion simulation markup language (used by several Tech-X codes)

21.7 Do you do data subsetting? Do you (or would you like to) make use of any database systems or other similar technology for performing subset extraction?

We use FastBit for performing data subsetting.

22. Analytics Applications

Code Analytic Apps (Q 22)

Python,
R,
VisIt,
several custom codes

23. Networking

Typically, network performance at the NERSC gateway is usually as good as it can possibly be (thanks!). Problems tend to lie elsewhere: e.g., connections to ORNL tend to show usable throughput of on the order

1-3Mb/s (as measured using iperf) even though they are connected directly to ESnet. This poor performance is not a NERSC issue; it is an ORNL issue.

24. Other HPC Support

A "startup-type" of allocation from NCCS/ORNL to provide support to stakeholders who compute/vis/analyze there.

25. Additional Information

26. Feedback and Project Requirements

We presently have a 1TB quota on NGF. When more space is required on a temporary basis, NERSC has been great about granting temporary quota increases. We foresee the need a larger NGF quota in the future due to an increasing level of activity.

27. Current Request Status

Not Finished: This request is not yet ready for review.

X
Finished: This request is finished and ready for review.